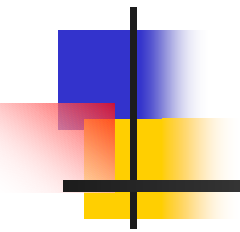

Lattice QCD

Project Status Report

April 25, 2006

Don Holmgren/Jim Simone

LQCD Project Drivers

- Delivering computing capability for lattice simulations
  - Stakeholders: FNAL theorists (Mackenzie, Simone, Eichten, Kronfeld, post docs), National LQCD collaboration (Bob Sugar, PI), FNAL CD, FNAL PPD, DOE HEP and NP
  - The Projects:
    - SciDAC Lattice Gauge Computing (2001-2006)
    - SciDAC-2 Lattice Gauge Computing (2006-2011)
    - “USQCD” (SC Lattice QCD Computing)

Scope

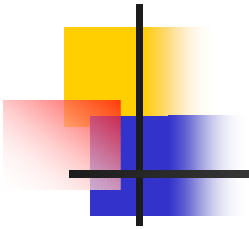

- SciDAC Lattice Gauge Computing Project (FNAL responsibilities in blue, some shared)
  - Optimization
  - Common libraries (QMP/QIO/QLA/QDP)
  - Runtime environment
  - Cluster prototyping
  - Physics production on cluster prototypes
  - Data interchange (International Lattice Data Grid)

Scope

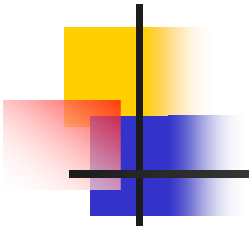

- SciDAC-2 Lattice Gauge Computing Project
  - Optimization for new hardware
    - Multicore
    - Opteron (or other NUMA)
    - Infiniband (or other network hardware)
  - Maintenance and improvement of libraries
  - Runtime Environment
  - Workflow Automation and Reliability of large systems
  - Algorithm development
  - Visualization
  - Hardware prototyping in support of USQCD
  - Data interchange (ILDG)

Scope

- USQCD (SC Lattice QCD Computing)
  - Operation of new and existing machines dedicated to LQCD at BNL, FNAL, JLab
  - Design and deployment of new machines
  - Project Management

Deliverables - SciDAC

- Optimizations
  - Scope is limited to SU(3) algebra (small complex matrix-vector multiplies)
  - Implementations:
    - Intel/AMD (SSE, SSE2, SSE3)
    - PowerPC (Altivec, native “double hammer”)
  - Essentially all of this work was completed years ago
  - Recent work
    - Maintenance for new compiler versions
    - The minimal changes for x86_64 ABI compatibility
    - Interactions with AMD/Cray to try and understand why performance is limited on Opterons
  - Status – Complete (more work in SciDAC-2)

Deliverables - SciDAC

- QIO library
  - This library abstracts hardware architecture details and allows I/O consistent with domain
    - Reads/writes SciDAC/ILDG file formats
    - Understands lattice layouts, parallel I/O
  - Role: Testing on clusters, and integration with dCache (libdcap)
  - Status:
    - Implemented in some physics codes (MILC via QDP, QDP++)
    - Libdcap issues with large files
    - Much work is needed to prepare for analysis of MILC fine lattices (8+ Gbyte propagator files) and super fine (17 Gbyte) in 2007

Deliverables - SciDAC

- QMP Library
  - Message passing library, essentially an LQCD-optimized subset of MPI, with multiple implementations which abstract details of underlying communications hardware
  - Role: Testing on clusters, native Infiniband implementation
  - Status:
    - Most physics production uses QMP over MPI (MILC, but not FermiQCD or Canopy)
    - Infiniband implementation has been delayed awaiting effort (called out by USQCD Project reviewers last May)
      - Deliverable for SciDAC-2

Deliverables - SciDAC

- Runtime Environment
  - Design and implementation of a user environment allowing portability between BNL, FNAL, JLab, and other sites for:
    1. Scripts and Makefiles
    2. Location of disk resources (scratch, HSM, parallel)
    3. Batch (submitting jobs, launching binaries)
    4. Data Grid (catalogs, utilities for moving files)
  - Status:
    - 1/2 implemented last summer, but decayed during OS upgrade in December
    - Collaboration has let this slip and is currently reviving
    - Progress promised by JLab/FNAL in next two months
    - Limited by effort at both sites

Deliverables - SciDAC

- Prototyping
  - Assessing performance and suitability of processors and networks for large scale clusters
  - Status (last 6 months)
    - Processor testing for JLab, FNAL USQCD clusters:
      - Dual core single Pentium (830D, 955EE)
      - Intel “Sossaman” (ultra low power – 110 Watt/dual box)
      - Intel “Dempsey” and “Woodcrest” (FBDIMM architecture)
      - AMD dual core Opteron
    - See USQCD review slides for results
    - Completed in time for JLab, FNAL procurements
    - Remote testing of Infinipath (flavor of Infiniband)

Deliverables - SciDAC

- Production prototype clusters
  - Decommissioned “w” (dual Xeon) in December
    - Helped transfer ¼ to accelerator simulation
  - Built “pion” in Spring/Fall
    - Largest to date: 520 single Pentium 640 nodes
    - Infiniband
  - Relocated “qcd” in December
    - To make room for New Muon Room 107 construction
    - 128-node Myrinet/Pentium 4E cluster
  - Operated “w”, “qcd”, “pion” (see USQCD)

Deliverables – SciDAC-2

- If funded, SciDAC-2 would start in July
- SSE:
  - Verify correctness in x86_64 mode (9/2006)
  - Optimizations for 64-bit mode (12/2006)
  - Revise for x86_64 ABI (6/2007)
  - Opteron-specific optimizations (9/2006)
- Native Infiniband QMP library
  - Report on prototyping (10/2006)
  - Implementation for “gen-1” (12/2006)
  - Implementation for “gen-2” (6/2007)

Deliverables – SciDAC-2

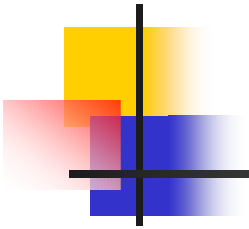

- Multicore support
  - OpenMP, thread, shm prototyping (12/2006)
  - Integrate with QLA/QDP (10/2007)
- Automated workflow
  - Specification for prototype complete (10/2006)
  - Profiling data generated for simulations (12/06)
- Monitoring and Control of Large Systems
  - Specification of prototype complete (9/2006)

Deliverables – SciDAC-2

- LQCD Runtime environment, GRID, ILDG
  - Deployment/testing of clients for metadata catalog access complete (12/2006)
  - Deployment/testing of clients for data movement complete (12/2006)
  - Integration with automated workflow system (starts Jan 2007)

Deliverables – SciDAC-2

- Hardware Prototyping – first year
  - Either Opteron “Socket F” or Intel Woodcrest
  - Infinipath flavor of Infiniband
  - Funds requested sufficient for ~ 16 nodes
  - Procurement would start as soon as funds are available
  - Need to provide results to USQCD project in time for their procurement (Nov 2006)

Deliverables - USQCD

- Project Management
  - OMB Exhibit 300 – BY08 by May 8
    - BY08 – first pass delivered April 6
    - BY07 – Delivered revisions in Aug 05, Dec 05
  - PM IT Certification
    - Level 1 Project Manager IT certification required
    - Holmgren – Finished training 12/05, certified
    - Banerjee – Finished training 4/06, certified (?)
  - WBS (cost and schedule tracking)
    - Maintained monthly by Bakul for BNL, FNAL, JLab
  - Performance (see milestones)
    - Maintained monthly by Bakul for BNL, FNAL, JLab

Deliverables - USQCD

- Monthly and quarterly reports
  - Informal reports are given to project and program managers (Jehanne Simon-Gillo, John Kogut, Sid Coon, Jeffrey Mandula) at phone cons monthly
  - Formal reports – DOE “Control Review Templates” – are submitted at the end of the 2nd month of each quarter
- External Reviews
  - “Integrated Baseline Review” – May 2005 at MIT
  - Project Progress Review – May 25-26, 2006 at FNAL

Deliverables - USQCD

- Deployment of new clusters:

Year	Deploy (Tflops)	Fermilab HW Budget
FY2006	2.0 (FNAL)	$1548K
FY2007	3.1 (JLab)	$25K
FY2008	4.2 (FNAL)	$1486K
FY2009	3.0 (FNAL)	$622K

- JLab “6N” cluster – 140/140 nodes (SciDAC/project)
  - Release to production May 1, 0.3/0.3 TFlops

Deliverables - USQCD

- Fermilab “Kaon” cluster
  - Approximately 450 dual processor nodes ($1.43M Project, $220K SciDAC/supplemental)
  - RFP released April 3, due back April 28
  - P.O. to issue about May 17 after DOE approvals
  - First rack by end of June, remaining racks by early August
  - Release to users by end of September
  - Had planned for 1.8 Tflops, but latest Intel/AMD performance will only give us 1.6 – 1.7 TFlops

Deliverables - USQCD

- Operation of existing clusters and QCDOC
  - Primary metric is delivered Tflops-yrs
  - Also, we track at each site uptime and response time on trouble tickets

Year	Deliver (Tflops-Yrs)
FY2006	6.2
FY2007	9
FY2008	12
FY2009	15

Deliverables - USQCD

- Progress through March:

Running Sums	FNAL	Jlab	BNL	Total	Pace	% Pace	Deficit
Oct	0.066	0.069	0.427	0.563	0.612	92%	0.049
Nov	0.120	0.124	0.766	1.010	1.087	93%	0.077
Dec	0.160	0.181	1.108	1.449	1.563	93%	0.114
Jan	0.224	0.237	1.458	1.919	2.038	94%	0.120
Feb	0.295	0.304	1.876	2.475	2.633	94%	0.158
Mar	0.377	0.355	2.226	2.958	3.108	95%	0.150

- Data in Tflops-yrs
- JLab “6N” will add 0.05 Tflops-yrs/month starting in May (and perhaps 0.025 Tflops-yrs in April)

Effort Profile

- Sources:
  - Holmgren, 50%+ PM, 50% other
  - Bakul Banerjee, 50%+ PM
  - Amitoj Singh, 100%
  - Kurt Ruthmansdorfer, 100%
  - Jim Simone, 50%
  - Bob Forster, 10-25%
  - Jim Kowalkowski, 25% (future unclear)
  - ESE Techs (as needed during construction)
  - Valery Sergeev (as needed during construction)
  - Total: 4+ FTE

Effort Profile

- Sinks
  - USQCD:
    - Administration/Hardware: 1.75 FTE
    - Site Management: 0.25 FTE
    - Software Support: 0.5 FTE (from base)
    - Project Management: 0.75 FTE (plus 0.5 from base)
  - SciDAC:
    - Cluster prototyping, software development: 2.5 - 3.0
  - Total: 6.25 – 6.75
  - This understates effort required. Help expected:
    - Pr. Sun (IIT sabbatical) will help on workflow 06/07
    - Interviewing candidates for #060075 in June
    - Still need more!

Risks

- Hardware
  - USQCD deploy goal requires 1.7 Tflops from 2006 cluster
  - New Intel architecture (FBDIMMs) has been late, slow, and hot
  - AMD Opteron was the fallback, and will likely win our RFP
  - Depending upon vendor pricing, we may fall short by 0.1 TFlops

Risks

- Facilities
  - All hardware is at New Muon
  - Health of New Muon affects availability and reliability of existing clusters
    - Room 108 uses chillers
  - New Muon Room 107 construction project
    - Beneficial occupancy slipped since August from April until June 15
    - June 15 looks solid
    - Room 107 should be available in time for first hardware at the end of June

Risks

- Infiniband Reliability
  - With “Kaon”, 95% of our capacity will depend on Infiniband hardware
  - Failures to date (500+ HCA’s, 33 leaf switches, 1 spine switch):
    - All of first 260 HCA’s replaced by vendor because of faulty soldering during manufacture
    - Since then, only 2 failed HCA’s
    - 3 failed leaf switches (bad power supplies)
  - We have strong relationship with Mellanox
    - Need to maintain this to assure hardware support

Risks

- Software
  - Linux stability
    - Very stable operations through November using a very old distribution (Fermi RedHat 7.3.x)
    - Upgraded to LTS 4.1 in December
    - Many kernel problems on head node through mid-February (crashes related to NFS)
    - Kernel.org kernel + patches stabilized head node
    - Currently battling NFS client stability on workers
      - 650+ clients do NFS mounts from lqcd.fnal.gov
      - Very low NFS traffic
      - NFS mounts are freezing – about to switch to UDP from TCP

Risks

- Software
  - dCache is our current parallel file system
    - Require arbitrarily large flat file system
    - Require large scalable I/O rates
    - Production throughput depends upon staging data products in and out of jobs
      - To dramatically increase in 2007
      - Data writing can overwhelm pool nodes because of independence of TCP/IP stack and dCache software
    - Mostly dccp usage now, but must move to libdcap
      - Large file support issue
      - Will move to x86_64 with “Kaon” – may reveal new issues

Risks

- Software
  - Infiniband software stack
    - We use “OpenIB”, development sponsored by tri-labs as part of ASCI
    - Very stable now, but various “nannies” are in-place to reset pieces (for example, the subnet manager)
    - We use two non-commercial MPI’s
      - Mvapich from OSU
      - Mpich-vmi from NCSA/UIUC
      - Both have required minor FNAL patching
    - We will have to move to OpenIB “Gen-2” once the ASCI community moves and “Gen-1” decays

Risks

- Software
  - Hints of a deep problem in OpenIB or MPI
    - By accident, MILC software with full message checking (check summing) was used in production starting in August
    - On some types of production, nearly deterministic checksum errors were found very late in runs
    - Onset of problem was eventually traced to a memory leak in MILC application code
    - Fixing the leak fixed the problem – but, the root cause was never determined

Risks

- Effort
  - We are running far too lean
    - Have not fulfilled ILDG responsibilities to my satisfaction since Eric Neilsen moved to EAG (12/2004)
    - Replacement opening has just been posted and interviewing should start in June
  - Key personnel – we are vulnerable to the loss of key personnel and should have succession planning
    - Very important to USQCD to overlap with JLab
    - JLab “6N” cluster very helpful to spread IB knowledge